Democratizing Machine Learning: Toward a Secure, Distributed, and Powerful Computer

Background

Machine learning (ML) has shown tremendous capabilities in various application domains over the past few years. An example of such an application domain is genomics, where ML can be applied to human gene expressions to predict certain disease therapy success rates. Thus, the application of ML in genomics has the potential to strongly improve cure rates while increasing health care cost-efficiency.

As single entities and software engineers often do not have the data, computing resources, or expertise to build their own ML models, Machine Learning as a Service (MLaaS) is emerging. MLaaS refers to an organization (i.e., the MLaaS provider) that offers to run an ML model on a cloud platform (e.g., Amazon Web Services, Microsoft Azure, Google Cloud). A client can then upload data to the cloud server, where the server then runs the MLaaS model with that client's data and returns the result to the client (e.g., a doctor uploads gene data to the cloud and, after an ML model runs there, it returns a result to the doctor of how successful certain therapies are predicted to be). However, several trust issues exist in this scenario, especially with sensitive data such as gene data. How can a doctor be certain that the MLaaS provider or the cloud company handles the data securely and does not abuse the data? How can the patient be sure of that? And how can the MLaaS provider be sure that the cloud company handles its models and its customers' data securely?

Trusted execution environments (TEEs) present themselves as a promising solution for those trust issues. A TEE is a secure area in a CPU, which performs computations while guaranteeing integrity and data confidentiality. In the described scenario, the cloud provider would, for example, sell TEE computing resources. The MLaaS provider then rents these resources and has a machine learning model ready for inference. The client (e.g., a doctor) can then run the service inside the TEE and use cryptographic techniques to ensure that nobody else can see the raw data or the result. Thus, the system provides stronger data confidentiality guarantees and has a higher chance of user acceptance.

Going even further, Blockchain technology can be used in the presented scenario to create new marketplaces for data, TEE resources, or machine learning models to be traded. Such a system has the potential to strongly increase health care quality while keeping cost low.

The goal of this research project is to implement ML algorithms in a TEE. Data from the health care and genomics domain can be used for it, however, the student may also propose other application domains. As this is a broad umbrella topic, a specific topic will be designed based on the student's interest and skills. Further areas where the student can focus on, are system security analysis, or blockchain-enabled data, resource, and model marketplaces. The ideal candidate has good programming skills as well as a solid understanding of hardware architectures and machine learning. Knowledge in TEEs and Blockchains is nice to have, but not necessary. The work allows you to gain deep knowledge and experience in rapidly growing fields: trusted hardware, machine learning, blockchain, and digital health.

Possible tasks

- Design, implementation, and evaluation of ML algorithms on a TEE (e.g., Neural Networks on Intel SGX)
- Design, implementation, and evaluation of blockchain-based marketplaces (e.g., for data, model, and TEE resource trading)
- Analysis of an integrated system incorporating security, privacy, and economic aspects

Introductory literature

Hynes, Nick; Dao, David; Yan, David; Cheng, Raymond; Song, Dawn (2018): A demonstration of sterling: a privacy-preserving data marketplace. In *Proceedings of the VLDB Endowment* 11 (12), pp. 2086–2089. Available online at http://www.vldb.org/pvldb/vol11/p2086-hynes.pdf.

Jones, Michael; Johnson, Matthew; Shervey, Mark; Dudley, Joel T.; Zimmerman, Noah (2019): Privacy-Preserving Methods for Feature Engineering Using Blockchain: Review, Evaluation, and Proof of Concept. In *Journal of medical Internet research* 21 (8), e13600. Available online at https://www.jmir.org/2019/8/e13600/pdf.

Kunkel, Roland; Le Quoc, Do; Gregor, Franz; Arnautov, Sergei; Bhatotia, Pramod; Fetzer, Christof (2019): TensorSCONE: A Secure TensorFlow Framework using Intel SGX. In *arXiv preprint arXiv:1902.04413*. Available online at https://arxiv.org/pdf/1902.04413.pdf.

Noah Johnson (Ed.) (2019): Building a Secure Data Market on Blockchain. Burlingame, CA: USENIX Association. Available online at https://www.usenix.org/conference/enigma2019/presentation/song.

Tramer, Florian; Boneh, Dan (2018): Slalom: Fast, verifiable and private execution of neural networks in trusted hardware. In *arXiv preprint arXiv:1806.03287*. Available online at https://arxiv.org/pdf/1806.03287.pdf.